SYSTEM AND METHOD FOR GEOREFERENCING DIGITAL RASTER MAPS

CROSS REFERENCE TO RELATED APPLICATIONS:

This application is a continuation of and claims priority from pending U.S. Application "System and Method for Georeferencing Digital Raster Maps" (09/537,849), filed March 29, 2000. Furthermore, this application is related to and claims priority from the following pending applications: "System and Method for Performing Flood Zone Certifications" (09/537,161), filed March 29, 2000 and "System and Method for Synchronizing Raster And Vector Map Images" (09/537,162), filed March 29, 2000 which are hereby incorporated by reference.

BACKGROUND OF THE INVENTION

Technical Field:

The present invention generally relates to geographic map processing and in particular to a system and method for associating digital raster maps with known geographic features. Still more particularly, the present invention relates to a system and method for associating specific points on digital raster maps with a geographic coordinate system.

Description of the Related Art:

A digital raster map is a computerized map image that resembles a conventional paper map in that it presents an image of the mapped area, but has no additional underlying data associated with the features of the map. A raster map is typically created by scanning a conventional paper map, and is a grid-based map composed of pixels (or dots) of color or black & white. Each pixel in the grid has can be referenced by the pixel coordinates, and has only one value, which indicates the color of that pixel. Raster images are commonly referred to as "bit mapped."

A vector map uses lines and polygons, rather than pixels, to describe an image. Unlike a raster map, in which the map image is simply stored as a grid of pixels, when a vector map is displayed, it is drawn from a set of underlying data. The vector map is created through the placements of nodes on a plane and connecting those nodes with lines. Vector lines can be attributed with tables of data such as elevations, values, names or other information relative to the line. Vector data can be displayed in three dimensions if the lines are attributed with z values, modified or changed relative to user need, or layered to allow for turning off and on the viewing of different information.

Because of their feature attribution properties, they are particularly useful for displaying geographic data. Vector maps are used to display boundaries or lines that denote the position and extent of features, such as county boundaries or lines denoting stream and river systems. It is also very easy to view or manipulate the data underlying a vector map, for example to view or change the elevation of a feature.

Also because vector maps are commonly drawn from a geographic data set describing the area shown, they are very easily, and generally inherently, georeferenced. Georeferencing is the process of relating source coordinates to referenced geographic coordinates., which are typically in standard latitude/longitude. An image or a vector file is georeferenced to be used within a mapping/geographic environment. In a vector map, the data from which the map is drawn will typically already include a geographic coordinate set.

Modern GIS systems normally make use of digital vector based map information. However, a vast legacy of paper-based map information exists. It is very expensive and time consuming to convert all of the information on these paper maps over to a digital vector format. In many cases the scope and expense of such conversions render them completely impractical. However, even when a complete conversion to digital vector-based format is not possible,

it is still possible to obtain some of the benefits of computerized map systems, first by converting the paper maps to digital raster maps (by scanning them), and then by georeferencing the raster image. After georeferencing, there should be a clear relationship between the pixel coordinates in the raster map, and the geographic coordinates of the feature represented by that pixel. It would be desirable to provide a system and method for georeferencing a raster map by associating points on that map with corresponding points on a previously-georeferenced vector map, or with reference points designated in the raster map which have known latitude and longitude.

5

SUMMARY OF THE INVENTION

It is therefore one object of the present invention to provide improved geographic map processing.

It is another object of the present invention to provide an improved system and method for associating digital raster maps with known geographic features.

It is yet another object of the present invention to provide an improved system and method for associating specific points on digital raster maps with a geographic coordinate system.

The foregoing objects are achieved as is now described. The preferred embodiment provides a system and method for georeferencing digital raster maps by associating specific points on the raster map with corresponding points on a georeferenced vector map or another georeferenced raster map. According to the preferred embodiment, a raster map and a corresponding vector map are simultaneously displayed to a user. The user then locates a common geographic point or feature on each map, and marks each of them as a unique point-pair. When the user has marked at least two point-pairs, the system then computes a georeferencing function, based on the pixel-coordinates of the points marked on the raster map and the corresponding geographic coordinates of the points on the vector map. Thereafter the geographic coordinates of any point on the map may be easily computed. The preferred embodiment provides that as more point-pairs are identified, the georeferencing function is modified for increased accuracy.

The above as well as additional objectives, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 depicts a data processing system in accordance with a preferred embodiment of the present invention;

Figure 2 is an image of an exemplary raster map, in accordance with the preferred embodiment;

Figure 3 is an image of an exemplary vector map, corresponding to the raster map of Figure 2, in accordance with a preferred embodiment of the present invention; and

Figure 4 is a flowchart of a process in accordance with a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

With reference now to the figures, and in particular with reference to **Figure 1**, a block diagram of a data processing system in which a preferred embodiment of the present invention may be implemented is depicted. Data processing system **100** includes processor **102** and associated L2 Cache **104**, which in the exemplary embodiment is connected in turn to a system bus **106**. System memory **108** is connected to system bus **106**, and may be read from and written to by processor **102**.

Also connected to system bus **106** is I/O bus bridge **110**. In the exemplary embodiment, data processing system **100** includes graphics adapter **118** connected to bus **106**, receiving user interface information for display **120**. Peripheral devices such as nonvolatile storage **114**, which may be a hard disk drive, and keyboard/pointing device **116**, which may include a conventional mouse, a trackball, or the like, are connected to I/O bus **112**.

The exemplary embodiment shown in **Figure 1** is provided solely for the purposes of explaining the invention and those skilled in the art will recognize that numerous variations are possible, both in form and function. For instance, data processing system **100** might also include a compact disk read only memory (CD-ROM) or digital video disk (DVD) drive, a sound card and audio speakers, and numerous other optional components. All such variations are believed to be within the spirit and scope of the present invention. Data processing system **100** is provided solely as an example for the purposes of explanation and is not intended to imply architectural limitations.

The preferred embodiment provides a system and method for georeferencing digital raster maps by associating specific points on the raster map with corresponding points on a georeferenced vector map or another georeferenced raster map. According to the preferred embodiment, a raster map and a corresponding vector map are simultaneously displayed to a user.

The user then locates a common geographic point or feature on each map, and marks each of them as a unique point pair. When the user has marked at least two point-pairs, the system then computes a georeferencing function, based on the pixel-coordinates of the points marked on the raster map and the corresponding geographic coordinates of the points on the vector map. Thereafter the geographic coordinates of any point on the map may be easily computed. The preferred embodiment provides that as more point-pairs are identified, the georeferencing function is modified for increased accuracy.

Figure 2 is an exemplary raster map, in accordance with the preferred embodiment. This exemplary map shows a scanned image from a Federal Emergency Management Agency (FEMA) paper map. This raster image shows a land area with flood zone indications, but would, in a computer system, contain no underlying data regarding the area shown.

Figure 3 is an exemplary vector map, corresponding to the raster map of Figure 2, in accordance with a preferred embodiment of the present invention. This map shows the same area as the map in **Figure 2**, but is created by a computer system from a database describing the locations of features such as the streets shown. Typically, each feature shown on a vector map such as this will already be georeferenced, in that the geographic coordinates of each feature will also be recorded in the underlying data

A digital map image is considered georeferenced if a pair of mathematical functions, f, and g, have been defined that can be used to convert back and forth between the coordinates of the map image (as defined by the pixels of the image) and the corresponding longitude and latitude of the location of that point. That is, f and g do the following:

1.   If (x, y) represents a location on the digital map image, then f(x, y) =(Lon, Lat) represents the longitude and latitude of the corresponding physical location.

2.   If (Lon, Lat) represents a physical location that lies within the region covered by the map, then g (Lon, Lat) = (x, y) represents the point on the digital map image that corresponds to that longitude and latitude.

Here, x and y represent the natural internal coordinate system of the map image. Typically, as described above, a digital raster map image uses the pixels of its image as a natural coordinate matrix. However, in most cases, a vector-based map image uses longitude and latitude as its internal coordinate system; if so, it can be considered to be trivially georeferenced already. Therefore, the functions f ( ) and   g ( )   above   are   non-trivial georeferencing functions required to convert back and forth between coordinate systems.

Following is a description of a georeferencing process in accordance with the preferred embodiment:

First, the raster map to be georeferenced (referred to as Map1) is displayed on a computer monitor in conjunction with another previously georeferenced map (referred to as Map2) **(step 400)**. Typically, Map2 will be a vector map, but could also be another raster map, for which georeferencing functions are already known. Map1 and Map2 can be independently manipulated for rotation, position and scale. The user will manipulate the maps until the region shown in Map1 is contained within the region shown by Map2.

Next, pairs of points representing identical geographic locations on Map1 and Map2 are identified, and the user uses a keyboard, mouse, or other pointing device to mark these point-pairs **(step 405)**. For each point-pair, one point is marked on Map1 and the other point is marked at the corresponding location on Map2. Examples of common georeferencing point-pairs include street intersections, places where rivers cross roads, mountain peaks, building locations, and other readily identifiable map features. If any points on Map1 have points of known longitude and latitude already marked on the map, then a

georeferencing point-pair can be obtained by marking that point on Map1 and moving to the point with identical latitude and longitude on Map2 and marking it as the other half of the point-pair. Alternatively, once the known points are marked with the known latitude and longitude, they can be used to determine the georeferencing functions without requiring a corresponding mark on Map2.

After a point-pair is marked, the x-y point on Map1 is assigned the corresponding latitude and longitude of its matching point on Map2 **(step 410)**.

When two georeferencing point-pairs are marked, the system will compute georeference functions for Map1 based on a linear transformation that allows an arbitrary rotation and predefined scaling, between the x-y domain and the latitude- longitude domain **(step 415)**.

To improve accuracy and allow increased generality, more georeferencing point-pairs will frequently be desired. To facilitate the process of gathering these additional georeferencing point-pairs, Map1 and Map2 are now selectively synchronized. This means that as either map is zoomed, panned, scrolled, or otherwise caused to display a different region, then the other map automatically does the same. This greatly increases the ease with which georeferencing point-pairs can be identified and marked on the maps, since when the user of the system locates an identifiable feature on one map, the other map will automatically be showing very nearly the same location on the other map.

Thereafter, when the user marks a georeferencing point on one of the maps, then the system automatically marks the corresponding point on the other map, based on its current georeferencing function **(step 420)**. The user can then make corrections to the location of a point marked by the system in this way **(step 425)**. The user may enter as many point-pairs as are desired; each additional point-pair will increase the accuracy of the georeferencing process.

When three or more georeferencing point-pairs are determined, the system computes georeference functions for Map1 based on a completely general linear transformation between the x-y domain and the latitude-longitude domain.

When four or more georeferencing point-pairs are determined, the general linear georeferencing functions are over-determined. This means that more than the required amount of information to compute the general linear georeferencing functions is available, but that it is not, in general, completely consistent. The system use the extra information contained in the additional georeferencing points to provide validation checks to protect against the possibility that some of the data points may be inaccurate (step 430). Points that deviate excessively with respect to a calculated standard error are presumed to be inaccurate and are omitted from the calculation of the georeferencing functions. Note that as new point are added, the system also rechecks points previously marked as inconsistent, to determine if those points should now be considered when recomputing the georeferencing functions.

The additional point-pairs allow the system to compute the general linear georeferencing functions which best fit the combination of all the available georeferencing point-pairs, so that the active georeferencing functions are revised with each new point-pair (step 435). According to the preferred embodiment, this is done by using a "least square" parameter fitting operation.

The user may then proceed o enter the next point-pair (step 440) . When the user is finished, the system stores the active georeferencing functions with the raster-map (step is 445). At this time, the raster map is considered fully georeferenced. When accessed at any future time, the system may simply retrieve the georeferencing functions, and apply them to find the latitude and longitude of any point on the raster map.

The process of determining he georeferencing function set from a set of

point-pairs is believed to be within the ability of one of ordinary skill in the art. The specific approach used by the system an method of the preferred embodiment is discussed below.

At any given point in the georeferencing process a set of points P, has been specified. Let $i \in P$ denote one of the points of P. Associated value x_i, y_i, Lat_i and Lon_i refer to the x and y coordinates of the point on the bitmap, and the latitude and longitude of the point on the earth's surface. We also define $A \subseteq P$ to be the subset of "active" points, i.e. points which are currently being used in the calculation of the georeferencing functions. Assume that we have selected a parametric family of functions F. Functions in this family map from (x, y) to (Lon, Lat). Any function $f \in calF$ has components $f_1(x, y) = Lon$, and $f_2(x, y) = Lat$. From within this family we seek the specific function $\hat{f}$, which (in a sense to be defined below) comes closest to making the following system of equations true:

$$\hat{f}(x_i, y_i) = (Lon_i, Lat_i) \quad for \ i \in A$$

Once determined, $\hat{f}$ will b the georeferencing function which is used to compute corresponding latitude and longitude values, (Lon, Lat) for any point, (x, y) on the bitmap. There are any number of possible ways to define the function that "comes closest to making (1) true." We shall follow a "least squares" approach also known in mathematics as an L_2 approach. This approach seeks to find the function, $\hat{f}$, which minimizes the sum of the squared differences between the actual and the predicted values of latitude and longitude. In other words, from among all the functions $f \in F$, $\hat{f}$ is the one which minimizes:

$$SSE = \Sigma\left[(f_1(x_i, y_i) - Lon_i)^2 + (f_2(x_i, y_i) - Lat_i)^2\right] \quad (2)$$

Among various alternative methods for choosing the function $\hat{f}$ are choosing it so t ht it minimizes the sum of absolute errors (rather than squared errors, or so that it minimizes the largest error. Other criteria are also possible.

In what follows we shall describe methods for computing georeferencing functions based on both a "rotational linear" parametric family of functions and a "general linear" family of functions. A knowledgeable practitioner will readily perceive that this approach may be easily generalized to other parametric families of functions. Obvious candidates include nonlinear parametric families of functions arising from the projection of the earth's surface onto a flat map.

The General Linear Case: In the general linear case, we let F be the set of all possible linear transformations which map from (x, y) to (Lon, Lat). Thus

$$\hat{f}(x,y) = \begin{bmatrix} \hat{a}_{11} & \hat{a}_{12} \\ \hat{a}_{21} & \hat{a}_{22} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} \hat{b}_{1} \\ \hat{b}_{2} \end{bmatrix} \tag{3}$$

for some choice of parameters $\hat{a}_{11}$, $\hat{a}_{12}$, $\hat{a}_{21}$, $\hat{a}_{22}$, $\hat{b}_{1}$, and $\hat{b}_{2}$. If the region covered by the map to be georeferenced is not too large, then this family of functions will contain a suitable function, $\hat{f}$ whose total error is quite small. In the case where the map to be georeferenced covers a larger area than this, then the curvature of the earth must be taken into account and F is not a suitable family of functions. In such a case, nonlinear functions must be used as mentioned above. We shall not pursue that case further, since it is a straightforward extension of the procedures used in the linear case.

To find $\hat{f}$ we seek the parameters which minimize

$$SSE = \sum_{i \in A} \left(a_{11}x_{i} + a_{12}y_{i} + b_{i} - Lon_{i} \right)^{2} + \left(a_{21}x_{i} + a_{22}y_{i} + b_{2} - Lat_{i} \right)^{2} \quad \text{The}$$

parameter values which minimize this expression are found by solving the following two independent systems of linear equations:

$$
\begin{bmatrix}
n & \sum_{i \in A} x_i & \sum_{i \in A} y_i \\
\sum_{i \in A} x_i & \sum_{i \in A} x_i^2 & \sum_{i \in A} x_i\, y_i \\
\sum_{i \in A} y_i & \sum_{i \in A} x_i y_i & \sum_{i \in A} y_i^2
\end{bmatrix}
\begin{bmatrix}
b_1 \\
a_{11} \\
a_{12}
\end{bmatrix}
=
\begin{bmatrix}
\sum_{i \in A} Lon_i \\
\sum_{i \in A} x_i\ Lon_i \\
\sum_{i \in A} y_i Lon_i
\end{bmatrix}
\qquad (5a)
$$

$$
\begin{bmatrix}
n & \sum_{i \in A} x_i & \sum_{i \in A} y_i \\
\sum_{i \in A} x_i & \sum_{i \in A} x_i^2 & \sum_{i \in A} x_i\, y_i \\
\sum_{i \in A} y_i & \sum_{i \in A} x_i y_i & \sum_{i \in A} y_i^2
\end{bmatrix}
\begin{bmatrix}
b_2 \\
a_{21} \\
a_{22}
\end{bmatrix}
=
\begin{bmatrix}
\sum_{i \in A} Lat_i \\
\sum_{i \in A} Lat_i \\
\sum_{i \in A} y_i Lat_i
\end{bmatrix}
\qquad (5b)
$$

These systems can be easily solve by well known methods, such as Gaussian Elimination, or LU factorization. The solutions yield the desired values of $\hat{a}_{11}$, $\hat{a}_{12}$, $\hat{a}_{21}$, $\hat{a}_{22}$, $\hat{b}_1$, and $\hat{b}_2$. It should be noted that equations (5a) and (5b) do not have a unique solution unless three or more non-colinear points are contained in A. Generally speaking, then, it requires 3 points to choose a georeferencing function from the family of general linear transformations. When there are four points or more, it is possible to compute a standard deviation of errors using the formula:

$$
s = \sqrt{\frac{\sum_{i \in A}\left[\left(\hat{a}_{11}x_i + \hat{a}_{12}y_i + \hat{b}_1 - Lon_i\right)^2 + \left(\hat{a}_{21}x_i + \hat{b}_2 - Lat_i\right)^2\right]}{n-3}}
\qquad (6)
$$

s is an estimator for the amount of error to be expected between actual and predicted latitude and longitude values.

Note that the inverse georeferencing function, $\hat{f}^{-1}$ that maps from (Lon, Lat) to (x,y) is readily obtained now, by inverting the function $\hat{f}$. Having done this, it is possible to compute a similar standard deviation of error for $\hat{f}^{-1}$ which is an estimate of error as measured in the bitmap coordinates.

The Rotational Linear Case: In the, so called, rotational linear case, we let *F* be the set of all possible linear transformations which map from (x,y) to (Lon,Lat) and which also allow for a known scale change caused by the difference in distance per degree of latitude and per degree of longitude at any point not on the equator. Thus .

$$\hat{f}(x,y) = \begin{bmatrix} \dfrac{\hat{\beta}_3}{\gamma} & -\hat{\beta}_4 \\ -\hat{\beta}_4 & -\hat{\beta}_3 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} \dfrac{\hat{\beta}_1}{\gamma} \\ \hat{\beta}_2 \end{bmatrix} \tag{7}$$

for some

choice of parameters $\hat{\beta}_1$, $\hat{\beta}_2$, $\hat{\beta}_3$, and $\hat{\beta}_4$, where $\gamma = \cos(Lat)$ is computed from the estimated latitude of the map to be georeferenced. Clearly, this is a special case of the more general linear transform described above, where we take

$$a_{11} = \beta_3 / \gamma, a_{12} = -\beta_4, a_{21} = -\beta_4,$$

$$a_{22} = -\beta_3, b_1 = \beta_1 / \gamma, b_2 = \beta_2.$$

In theory, if the region covered by the map to be georeferenced is not too large, then this family of functions will contain a suitable function,

$\hat{f}$ whose total error is quite small. In the case where the map to be georeferenced covers a larger area than this, then the curvature of the earth must be taken into account and F is not a suitable family of functions. In such a case, nonlinear functions must be used as noted above. In practice, the general linear transform will yield results which are somewhat superior to the rotational linear transform due to the former's innate ability to use arbitrary scaling factors. These arbitrary scaling factors might compensate, for example, for scaling flaws in the digital scanner used to scan in the maps. The primary advantage of the rotational linear transform is that it can be computed with only two reference points (unlike the general linear transform which requires three reference points).

To find $\hat{f}$ we seek, as before, the parameters which minimize SSE as defined in equation (4).

The parameter values which minimize this expression are found by solving the following system of linear equations:

$$\begin{bmatrix} n & 0 & \sum x_i & -\sum y_i \\ 0 & n & -\sum y_i & -\sum x_i \\ \sum_{i \in A} x_i & -\sum_{i \in A} y_i & \sum_{i \in A}\left(x_i^2 + y_i^2\right) & 0 \\ -\sum_{i \in A} y_i & -\sum_{i \in A} x_i & 0 & \sum_{i \in A}\left(x_i^2 + y_i^2\right) \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{bmatrix} =$$

$$\begin{bmatrix} \gamma \sum_{i \in A} Lon_i \\ \sum_{i \in A} Lat_i \\ \gamma \sum_{i \in A} x_i Lon_i - \sum_{i \in A} y_i Lat_i \\ -\gamma \sum_{i \in A} y_i Lon_i - \sum_{i \in A} x_i Lat_i \end{bmatrix}$$

These systems can easily solved by well known methods, such as Gaussian Elimination, or LU factorization. The solutions yield the desire values of $\hat{\beta}_1$, $\hat{\beta}_1$, $\hat{\beta}_1$, and $\hat{\beta}_1$ which in turn yield the desired values for $\hat{a}_{11}$, $\hat{a}_{12}$, $\hat{a}_{21}$, $\hat{a}_{22}$, $\hat{b}_1$, and $\hat{b}_2$.

5

It should be noted that equation (8) does not have a unique solution unless two or m re points are contained in A. Generally speaking, then it requires two points to determine a georeferencing function from the family of rotational linear transformations. When there are three points or more, it is

10   possible to compute a standard deviation of error, s using the formula:

$$s = \sqrt{\frac{\sum_{i \in A}\left[\left(\hat{a}_{11}x_i + \hat{a}_{12}y_i + \hat{b}_1 - Lon_i\right)^2 + \left(\hat{a}_{21}x_i + \hat{a}_{22}y_i + \hat{b}_2 - Lat_i\right)^2\right]}{n-2}}$$

s is an estimator for the amount of error to be expected between actual and predicted latitude and longitude values.

Note that the inverse georeferencing function $\hat{f}^{-1}$ that maps from

15   (Lon, Lat) to (x, y) is readily obtained now, by inverting the function $\hat{f}$. Having done this, it is possible to compute a similar standard deviation of error for $\hat{f}^{-1}$ which is an estimate of error as measured in the bitmap coordinates.

Automatic Error Detection and Handling When individual points are being assigned x, y, Lon, and Lat values, there is always a potential for

20   error. To reduce the risk of incorrect georeferencing resulting from such errors, certain error handling procedures are built into the georeferencing process. The fundamental conk t is that of detecting a "bad" point and then removing it from the set of active points, A. Note that removing point from A will not delete the information associate, with that point, but it will cause the georeferencing

parameters to be completely uninfluenced by that point. We o not wish to remove the point entirely, since it may be determined at a later stage of the georeferencing, that the point was not really bad at all, and should be-used in the georeferencing calculation. This will be clarified shortly.

5

Detecting Bad Points The following steps outline the bad point detection process using the general linear transform approach to georeferencing.

1.   Begin by placing all existing points into the active set, A

10   2.   If there are fewer than five active points then you are done,

Otherwise, for each of the currently active points in turn, move it   (call it point k for the sake of convenience) temporarily out of   the active set, and then calculate the resulting inverse   georeferencing function (call it $\hat{g}^{(k)}$) and its corresponding SSE_k.   Also, calculate the difference between   the predicted value and the   actual value $\delta_k = \left| \hat{g}^{(k)}(Lon_k, Lat_k) - (x_k, y_k) \right|$.

15

Make a note of the values, δ_k and δ_k/SSE_k.   Return point k to the active set (and move on to the next value of k.

3.   From among the results found in step 2 above, find the point, k, with the largest value of $\delta_k/SSE_k \rangle c_1$ and δ_k which also

20   satisfies $\delta_k/SSE_k \rangle c_1$ and $\delta_k \rangle c_2$, where c_1, and c_2 are some constants which are set according to the general level of accuracy to be expected on the particular type of map which is being georeferenced, the current number of active points, and the dots per inch of the scanned image. If there is such a point then mark it as bad (by removing it from the active set) and return to step 2 above.

Otherwise you are done.

There are several things to note about this procedure. One is that it allowing the value of c_1 and c_2 to change with the number of active points, makes it possible for the georeferencing system and method to utilize points which it might originally determine bad or inconsistent after a large enough sample of points has been gathered to make it clear that a lesser level of accuracy is all that can be achieved on this map. Another observation is that using this procedure it is impossible to reduce the number of active points down to less than four (unless you started with less than 4 in which case this procedure does not apply at all). This scheme means that as each new point is added, all points determined so far are considered, even those which had previously been marked bad. Thus early "misjudgments" on the part of the system can be corrected later, in light of new point information.

The same bad point detection process, can also be implemented using the rotational linear transform approach. In this case the method is capable of reducing the number of active points down as low as three (rather than four for the general linear transform approach outlined above). This can be useful when dealing with small sets of active points.

When the system is georeferencing raster map images that cover a large enough area so that the (nonlinear) curvature of the earth is a source of significant error, the system can calculate nonlinear georeferencing functions whose form corresponds to the map projection that was used to create Map1.
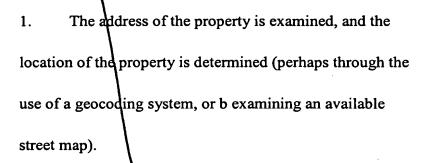
A specific example of the operation and application of the preferred georeferencing method may be shown with reference to the "Flood Zone Determination" business. The Federal Emergency Management Agency (FEMA) publishes a library of tens of thousands of paper maps showing various types of flood zones and their locations in the United States. A flood zone determination on a property is frequently done in the following way:

1. The address of the property is examined, and the location of the property is determined (perhaps through the use of a geocoding system, or b examining an available street map).

2. A map analyst attempts to determine which of the many thousands of FEMA flood maps will contain this property.

3. The map analyst goes to a map storage area and retrieves the desired map, often examining several maps before making a final selection.

4. Having retrieved the paper map, the map analyst next determines where, precisely, the property is located on the map.

5. Finally, the map analyst examines flood zone notations on the map at the property's location in order to determine its flood-zone status.

When performed using paper maps, the above process is difficult and quite time consuming . A database of scanned raster map images, alone, can be used to reduce the time and effort expended on step 3 above. However,

georeferenced raster map images can also be used to eliminate virtually all work from steps 2 and 4 above. Since those are generally the most time consuming steps under the current process, the value of georeferenced flood maps is considerable.

Using georeferenced flood map raster images, steps 2 and 4 above, are replaced by:

2.   A computer system combines the pre-designated outlines of the raster map and the georeferencing information to obtain a polygon expressed in terms of latitude and longitude that outlines the region included in each flood map. Then the system determines which of the polygons contain the address in question, which is done using a "point-in-polygon" algorithm. At the conclusion of this process, the computer system has identified a map panel (or perhaps a small number of map panels) that contains the address.

4.   Since the latitude and longitude of the property are known (by virtue of a geocoding phase), the computer system can use the georeferencing of the map panels to locate the property on each of the panels found above, thus largely eliminating any need for he map analyst to scan the flood map for the address location

Therefore, it is easily seen that by using georeferenced raster maps instead of paper maps, the process of determining flood zones is substantially automated and much more efficient.

It is important to note that while the present invention has been described in the context of a fully functional data processing system and/or network, those skilled in the art will appreciate that the mechanism of the present invention is capable of being distributed in the form of a computer usable medium of instructions in a variety of forms, and that the present

invention applies equally regardless of the particular type of signal bearing medium used to actually carry out the distribution. Examples of computer usable mediums include: nonvolatile, hard-coded type mediums such as read only memories (ROMs) or erasable, electrically programmable read only memories (EEPROMs), recordable type mediums such as floppy disks, hard disk drives and CD-ROMs, an transmission type mediums such as digital and analog communication links.

While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention.